AAS 93-031

Low Cost Spacecraft Computers: Oxymoron or Future Trend?

Robert M. Manning
Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

Over the last few deeades, application of current terrestrial computer technology in embedded spacecraft control systems has been expensive and wrought with many technical challenges. These challenges have centered on overcoming the extreme environmental constraints (protons, neutrons, gamma radiation, cosmic rays, temperature, vibration, etc.) that often preclude direct use of commercial off-the-shelf computer technology. Reliability, fault tolerance and power have also greatly constrained the selection of spacecraft control system computers. More recently, new constraints arc being felt, cost and mass in particular, that have again narrowed the degrees of freedom spacecraft designers once enjoyed.

This paper discusses these challenges, how they were previously overcome, how future trends in commercial computer technology will simplify (or hinder) selection of computer technology for spacecraft control applications, and what spacecraft electronic system designers can do now to circumvent them.

INTRODUCTION

A ncw effort is afoot in the scientific space community to define a myriad of ncw space science research objectives centered on the concept of very small, simple, lightweight, and low cost spacecraft that will widen the scope and number of space exploration objectives and generally enhance space science accessibility. A cornerstone of this concept lies in the perception that the spectacular advances in terrestrial electronics technology (particularly in computer, actuator, and sensor technology) pave the way for similar advancements in spacecraft electronics technology. In particular, fcw would argue that the performance / price, performance / mass, and performance/ power ratios for terrestrial computer systems have steadfastly remained on exponentially increasing curves for at least the last two decades.

However, in the traditional world of large, expensive, unique, long life spacecraft, the change in these figures of merit, although noteworthy, have not been as dramatic. Despite the high marginal cost per kilogram [1] and cost per watt of traditional spacecraft payloads, spacecraft computers are often characterized as having poor performance with respect to these key cost drivers. This is largely because traditional spacecraft designs have employed redundant, high reliability, radiation hardened electronics to greatly enhance the probability of mission success. This insurance policy comes with a high price tag: increased recurring component costs, increased non-recurring component costs, greatly decreased component selection, increased mass and power. Since the mission costs were already high (and the cost of failure proportionately higher), these costs have been considered acceptable.

On top of all these costs, the complexity of these spacecraft requires that the flight computer and its associated input/output (1/0) electronics bc procured well in advance of mission launch so that exhaustive system-level functional and environmental testing may proceed. Also, to mitigate risk even further, the mission designers may choose to restrict themselves to components that have shown a certain degree of maturity in the marketplace. These all add to the well-known 8-year (or more) technology gap [2] between terrestrial

computer technology and launch-day space computer technology, as well as adding to the high costs that have perennially haunted the spacecraft computer industry.

In order to meet the demands of the military, the 1980's saw a renaissance in funding for space electronics. Programs such as NASA's on-board computing technology programs [3], DoD's VHSIC (Very High Speed Integrated Circuit) program [4], and the Airforce Space Division's various space-hardened microelectronics programs (GVSC, ASCM, and RH-32) all have made, and are continuing to make, contributions to the advancement of flight computing. Today's traditional spacecraft developers rely heavily on the technologies developed under these programs. JPL's Cassini spacecraft to Saturn directly inherited computer and memory component technology developed under some of these programs (i.e., 256 Kbit SRAM, and IBM's GVSC 1750A). However, increased funding for such programs is diminishing. Riding on the coat-tails of government-sponsored defense space electronics research and development may not be as common in the 1990s.

The new demands made by the very small, low cost spacecraft concept invite a re-evaluation of the problems and solutions that traditional spacecraft designers have faced; to step back and ask if the impressive gains made in the terrestrial computer market might offer hope of making inroads against the high costs associated with traditional spacecraft computer electronics, and to sec if the new spacecraft designers themselves can take steps to mitigate these costs.

SUMMARY OF SPACE-UNIQUE ISSUES

Many of the problems that are unique to space-based computer systems many not be solved for the designers of small spacecraft, but they arc certainly well understood [1] [5] [6] [7] [8]. Total integrated dose radiation effects, being among the most nagging of them, has held the interest of the defense electronics industry since the early 1960's. However, it was only in the late 1970's, with the advent of small device feature sizes (and an active solar climate) that general interest in the non-destructive effects of highl y energetic particles has taken hold [9]. Likewise, the desire for low mass and low power computers has a been an ongoing concern, but for the most part this desire has taken a back seat to the greater need for reliability.

Performance

Although moving at a measured pace, spacecraft computer performance has by no means been static. As is evident in Fig. 1, with the exception of high power bipolar discrete component computers, spacecraft computer performance has tracked about an order-of-magnitude behind the currently available terrestrial computer capability. Computer performance has generally been quantified in units of MIPS (Millions of Instructions Per Second). Although a widely used indicator of computer (processor) performance, there is no single standard benchmark for comparing the MIPS rating of one machine with another although some have gained popularity [10].

An important point to recall is that most sequential processors can be classified by their instruction sets. Until recently, most of the microprocessors that have been in general usc (MC680XX, i80x86, NSC320xx, etc.) and virtually all of the space computers used to-date (CDC's 469, Litton's 4516E, RCA's SCP-234 & 050, Teledyne's Meca 43, Rockwell's DF-224, Delco's M362S, IBM's NSSC-I and II, Applied Technology's ATAC-16MS, Sandia's SA2689, and everybody's MIL-STD- 1750A, etc.), have had rich, complex instruction sets for the assembly language programmer or the high level language compiler writer to draw upon. This class of computer is known as CISC or Complex Instruction Set Computers. The instruction set complexity generally meant that less instructions were required to do a given task. This allowed relatively slower (and smaller) memory to keep the processor operating at peak performance. Today, most ncw high performance

microprocessors (e.g., SUN SPARC, MIPS R6000, RS6000, i8800, etc.) rely on small well-tuned instruction sets. This class of computer is widely known as RISC or Reduced Instruction Set Computers. By constraining the instruction set, RISC processors [11] use multi-stage pipelined (i.e., parallel) instruction processing, effectively allowing the compiler writer to organize the flow of instructions so that the pipeline can be more optimally filled. This has had several implications: much higher processing throughput potential, an increase in the number of instructions required to do a given task, many more instruction accesses to memory (partially offset by instruction caches), and making the MIPS rating of a RISC about a factor of 2 or 3 times higher for a given task than a CISC [12]. Therefore normalization of the processor's MIPS rating for comparison purposes is called for when comparing these machines.

The intrinsic relationship that average memory access time has on RISC throughput means that space applications that use RISC must pay close attention to the memory architecture. Flight RISC performance can bc greatly enhanced if an associated space-qualified cache is included, and even then, low access time memory chips may be nccessary.

Many of the current space processor developments center on using radiation hardened and/or fault tolerant versions of terrestrial commercially popular 32-bit RISC processors. Examples of 32-bit RISC developments underway in the US include: RH-32 (Honeywell and TRW/McDonald Douglas), Rad6000 (IBM), and R3000 (LSI Logic, Harris Semiconductor), RTX 2000 (Harris Semiconductor). In Europe, ESA [13] is also supporting development and qualification testing of versions of SPARC (Phillips), MIPS R3000, (Siemans), ARM 2 (Advanced RISC Machines), and the T800 Transputer (SGS-

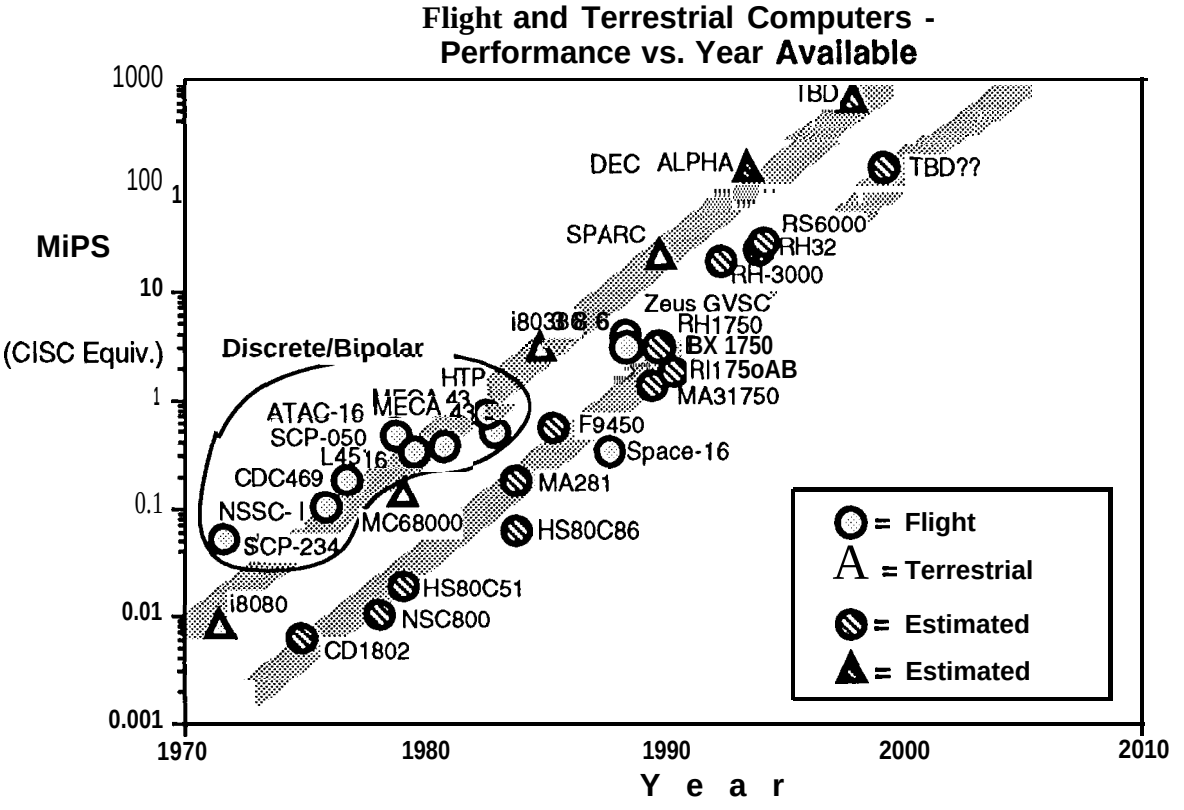

Figure 1

Thomson).

By far the most common single spacecraft computer instruction set is the Mil-Std-1750A. A large number of vendors provide class S (or B) versions of this 16-bit machine with various radiation hardening and performance characteristics (Allied Signal, CDC, Honeywell , IBM, LSI Logic, Pace, Rockwell, UTMC, GEC-Plessey, Westinghouse, to name a few) [14].

It can be argued that with the large number of suppliers, sales competition will (eventually) make 1750 computers the low cost leader in small spacecraft applications. This is only slowly beginning to happen. Because the primary user of 1750 flight computers has been the military, only the standard government-approved high level languages (Ada and Jovial) have received military support. This contrasts sharply with the trends in the commercial world to usc languages like C, C++, and Forth, as well as off-the-shelf embedded operating systems like VRTX, Wind River, SC/Forth, and Chorus. As a result, non-military applications have been reluctant to abandon these for the more complex (some would say antiquated) and expensive military languages and operating systems. The outcome isn't clear, but either the non-military low cost spacecraft designers will adopt 1750 and the military languages, use of non-military languages and operating systems for the 1750 will become more common, or low cost spacecraft designers will abandon the potential hardware cost savings and opt for trying other non- 1750 options such as the emerging radiation hardened 32-bit processors or even custom approaches.

The latest silicon compiler VLSI and system design tools (e.g. Mentor, Cadence, Synopsis, i-Logix, etc.) have enabled small teams of designers [15] to develop custom flight computers from the ground up at greatly reduced cost over previous efforts [16]. These efforts show much promise in tailoring the hardware to the application using inexpensive commercial processes and porting the design to rad-hard foundries. If affordable, this approach may very well pave the way for greatly reduced mass, power, and volume of future lightweight spacecraft.

Power, Mass, and Volume

An interesting figure of merit for all computers (and especially the small low power spacecraft computer) is the number of MIPS pcr watt of power expended. Surprisingly enough, in the late 1970's, space computer technology was actually very c1osc to the current commercial technology using this measure. This is because terrestrial computers generally used fast, high power bipolar chip technology (≈1 MIPS at 250 W or 0.004 MIPS/W). Spacecraft, on the other hand, were among the first applications of the slower, but much lower power, CMOS computer technology (0.01 MIPS at less than 5 W or less than 0.002 MIPS/W). As the advantages of CMOS technology have caught hold, terrestrial technology has rapidly outpaced space technology in this metric as well. Figure 2 shows a plot of maximum performance as a function of power of some past and present flight computers and where they stand with respect to performance power ratios [17] **[18]** [19].

As can bc seen in Fig. 2, the device technology greatly affects the MIPS/Watt ratio. This is because the power expended is directly related to clock rate, average feature size, the degree of parallelization, and chip transistor technology. The energy dissipated can be expressed in the following approximate relation:

$$P = CV^2f + VI_q \qquad (1)$$

Where C is the "total" load capacitance, V is the supply voltage (typically 5 volts), f is the clock frequency, and Iq is the computer quiescent current.

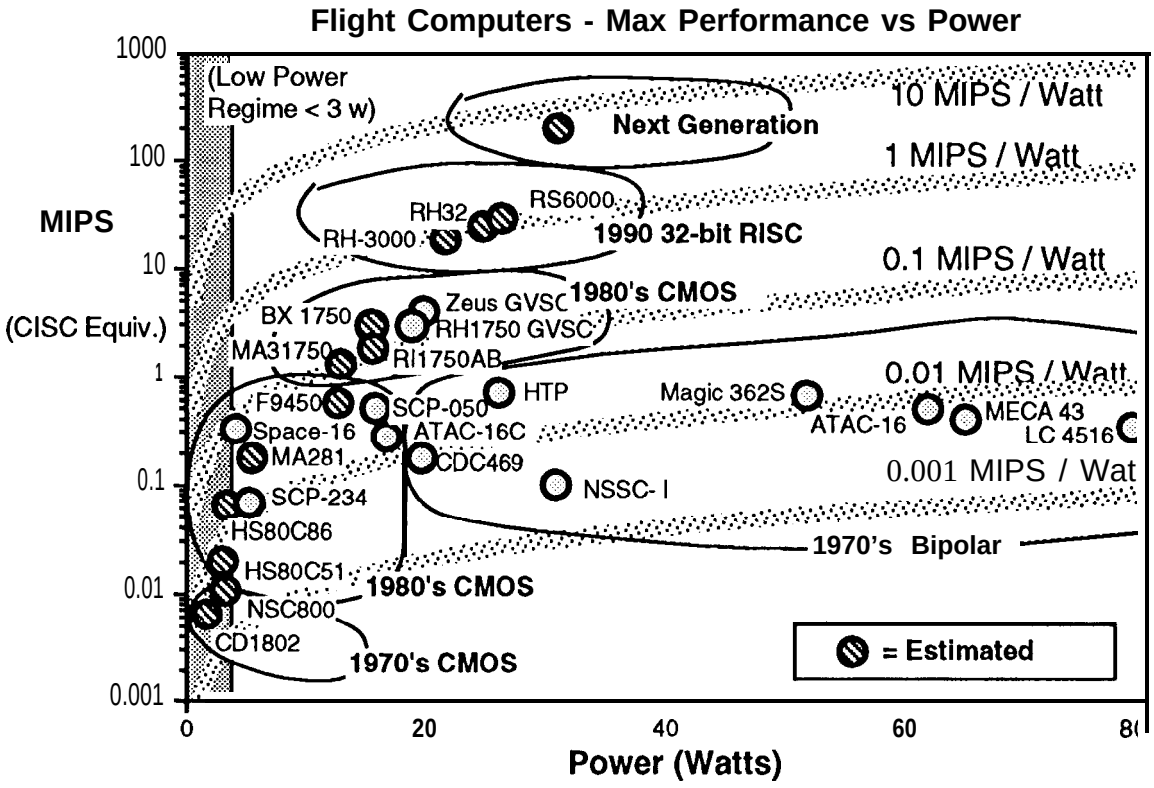

Figure 2

Since a single quantity for load capacitance and frequency can't easily be calculated for a whole computer, Eq. 1 must be separately applied to each switching circuit throughout the computer and summed. However, at the chip level, to first order, the load capacitance is proportional to the minimum feature size and frequency is proportional to the computer's input clock frequency (typically the spacecraft designer's only degree of freedom). This means that, all else being equal, the power drops as the computer manufacture scales down the size of the component technology.

For some of the computers shown in Figures 1 and 2, comparable performance was estimated and scaled to CISC MIPS based on published performance figures for those components. Where not available, power was estimated based on board-level power utilization of comparable computers with similar frequency, part count, I/O structure, and chip technology. Power associated with board I/O was assumed to be zero. For all of the computers, the indicated performance for specific computers is accurate to no better than a factor of two.

For a fully CMOS computer I_q is very small (even on spacecraft, some components are not all CMOS; typically the 1/0 circuits contain some low impedance bipolar circuits), therefore, the power dissipated in CMOS computers tends to drop to near zero when the computer's oscillator frequency is dropped to zero (see Fig. 3).

The computer's input clock frequency is also directly proportional to the processor performance (MIPS). In fully CMOS computers, doubling the clock frequency doubles the performance. This feature has great utility for the spacecraft designer, and has been used in applications requiring warm standby computer operation (such as for warm battery-backup sparing), and power-sharing during peak power loads. In fact, scaling the computer clock frequency down to the absolute minimum performance requirement to stay within the

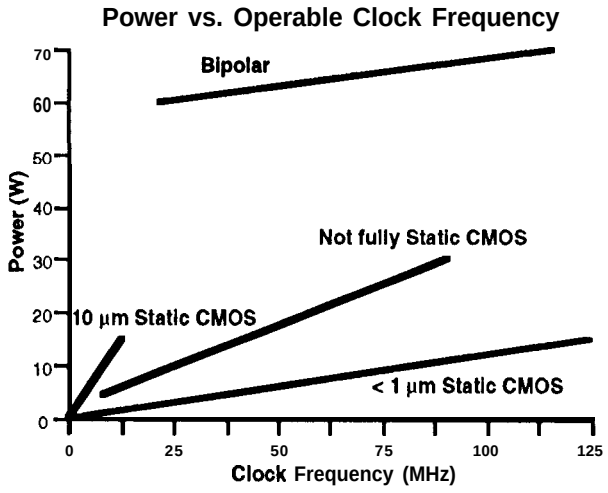

Figure 3

power budget is now common.

This advantage makes for a subtle case in favor of the use of state-of-the-art computer technology for small spacecraft, even in applications that do not require high performance. As discussed above, in nearly all spacecraft designs, it is highly desirable to have a high performance to power ratio. Since computers with the highest ratio also have the lowest load capacitance, and load capacitance scales with minimum feature size, it follows that the computers with the smallest geometry (i.e., state-of-the-art) will provide the highest performance to power ratio. To stay within the power limitations, one simply reduces the systcm clock frequency until the power is within the design margin.

Note from figure 2 that the highest performing processors (with the smallest feature sizes) tend also to reside on the highest performance to power ratio curve. Although there arc other factors (such as electron mobility and computer architecture), this is because the maximum performance a processor can have is limited by the feature size in two ways: the maximum rate the signal can transition between logic levels is governed by the load capacitance, and the signal propagation time which is governed by wire length, both are functions of feature size. Power and current density play a role in limiting the maximum performance, however these too arc mitigated by usc of CMOS and reduced feature size.

From Eq. 1, a reduction of the supply voltage would seem a good method for greatly reducing the power. Efforts are underway [20] [21] to do just that, however the electronics industry has not strayed far from 5 volt (CMOS) logic for the past decade or two, primarily because of intrinsic circuit parameters (device thresholds arc typically around a volt or two) and the desire to maintain noise immunity. In high speed computer designs packed with high impedance, 200 MHz bandwidth digital signals, cross-talk, inductive ground bounce, and transmission line effects make noise margin an important commodity.

The trend in terrestrial computer technology toward higher performance, and therefore smaller feature size, presents another benefit to the small spacecraft designer: reduced mass and volume. New technology such as multi-chip modules (MCMs), 3-D packaging and very high density connectors and backplanes show high promise on microspacecraft.

The trends in terrestrial computer technology seem to be naturally adapted to future small, low mass, power and volume spacecraft computer systems. The next sections will explore the issues that complicate that technology transfer.

Total Dose Radiation Effects

There are two primary radiation-induced failure phenomena that plague the design of space electronics: Total Ionizing Dose (TID) damage and Single Event Effects (SEE). Together they represent the single most expensive and limiting aspect of spacecraft computer design.

Total dose radiation damage is caused by high energy photons (Gamma or X-rays) or high energy particles (e.g. neutrons, electrons). Both have similar effects on the electronic properties of semiconductor materials. In onc kind of damage, atoms in the lattice are displaced, causing net changes in the semiconductors properties. Annealing sometimes

reverses this phenomenon, but over time degradation inevitably occurs. In another, electron-hole pairs are created. Recombination usually makes this a transitory phenomena in the bulk semiconductor, but when this occurs in insulating material, reeombination may never occur. Instead these pairs may distribute themselves near the conductive channel in the gate oxide in such a manner that they effectively change the MOS transistor's threshold voltage.

The TID radiation threshold for a component is expressed in units of rads (radiation absorbed dose equaling 100 ergs / gm) for photons and fluence (in total number of neutrons / cm^2) for neutron dose. It is the maximum amount of energy required to be deposited on that component material (typically silicon) before parametric damage occurs.

Typically, the required component radiation threshold is set between 2 and 10 times the dose expected over the mission life. This safety factor is known as the radiation design margin (RDM) and is placed on the mission design by the contracting agency or the projeet manager. The RDM is as large as it is primarily because the actual dose the spacecraft will receive may vary by up to an order of magnitude duc to solar fluctuations [22].

Several standard radiation hardness levels for components can be found in Mil-M-38510. Radiation hardness levels in 38510 for either class S or B components vary from none ("/") to 1 M rad (Si) TID and 2x1012 n/cm^2 neutron fluence ("H"). Standard radiation dose qualification tests are specified in Mil-Std-883 methods 1017 (neutrons) and 1019 (gamma). For new computer component designs, a fcw electronic components are irradiated until parametric shifts and subsequent failure occurs. If the process line is a commercial line that is not expressly designed for high radiation dose tolerance, lot to lot variation of the TID tolerance may also require radiation testing of samples from each lot (with associated recurring costs that may exceed $1500 / lot).

Since radiation is affected by the geomagnetic shielding of the Earth's (or other) magnetic field, the actual radiation dose rates vary considerably depending on orbit and orbital inclination. The following table shows some approximate expected doses as a function of orbit. Note that these will also vary as a function of solar output (see below).

Orbit	Inclination	Dose Rate
Low Earth (200 -1000 km)	<28°	100- 1k rad (Si) / year
Low Earth (200 -1000 km)	>28°	1k - 10 k rad (Si)/year
Medium Earth (1000 -4000 km)	any	100 k rad (Si) / year (Van Allen)
High Earth (≈36 k km)	any	> 10 k rad / year
Interplanetary	n/a	5 k - 10 k rad / year

Some chip fabrication processes are expressly designed for radiation dose tolerance. This entails subtle processing changes that bias the transistor's thresholds away from the threshold shifts encountered after radiation. Another technique involves making the gate oxide very thin so that oxide volume is reduced to minimize the effect of trapped charges on gate thresholds. These processes tend not to greatly influence the chip's design, therefore it is possible (provided the design is compatible with the layout design rules for that process) to take an existing chip that was designed for non-radiation hardened processes to be "ported" to a radiation hardened process with little or no modification.

As chip geometries shrink, these same process-related properties arc gradually appearing as standard features in some process lines that were not specifically intended for production of radiation hardened components. Some manufacturers have found that their processes yield reasonably high TID thresholds (e.g. 5 k - 10 k rad). However, unless

controlled, radiation tolerances may vary significantly from lot to lot from normal process variation or as a result of process enhancements [23].

Chip designs that can be easily transferred from commercial process lines to radiation hardened process lines can greatly reduce costs. Efforts at porting designs have enabled some groups [23] to develop semi- or full custom ASIC designs on inexpensive processes (such as provided by MOSIS) and later pay fabrication costs for the more expensive radiation hardened lines only after the design has been functionally tested. Foundries providing radiation hardened processes include Honeywell's RICMOS process (both full and semi-custom ASICs), and UTMC'S and LSI Logic's gate arrays. Radiation hardened fabrication costs for ported designs range from tens of thousands for some gate arrays to well over $100 k for full-custom class S designs.

Non-radiation dose hardened components can be used if adequate shielding is present. High density (high atomic number) shielding materials such as tantalum (often in the form of spot shielding) provide highest radiation absorption at the expense of mass. This shielding is typically placed adjacent to the components to minimize mass. Care must be taken to avoid the effects of secondary emissions (e.g. bremstrahlung radiation) from the shielding material itself. These can be minimized by using low atomic number shielding materials (e.g. aluminum) on the exterior of the computer [6].

As computer volume continues to shrink, it may become cost effective to usc very low volume, relatively inexpensive, highly integrated computer designs that are not particularly hard to high accumulated radiation doses (i.e., 1 k -5 k rad) and instead pay the cost of increased shielding mass. This approach is being taken in the design of JPL's solid state recorder (SSR) for the Cassini mission. The class B dynamic random access memories (DRAMs) in the SSR arc not radiation hard, so the SSR will be totally surrounded with at lcast 1 cm of aluminum shielding.

When making comparisons of the mass of various flight computers, it is important that the radiation tolerances of the machines meet or exceed the mission's radiation requirements (see below). This is because radiation shielding mass may account for a significant fraction of a flight computer's total mass. Whereas a flight computer made up entirely of radiation hardened components may require very little or no radiation shielding.

Single Event Effects

Unfortunately, TID-hardened processes do not necessarily ensure hardness to another class of failure mechanisms commonly known as Single Event Effects (SEE) or Single Event Phenomena (SEP). For digital microcircuits this failure mechanism can be further broken down into two primary effects: Single Event Upset (SEU) and Single Event Latchup (SEL). The former tends strictly to be a (physically) non-destructive phenomenon that injects noise into component circuits causing unwanted bi-stable device state transitions (bit upsets). The latter is quite often destructive to the components that exhibit it. SEL is induced when a "wake" of electron-hole pairs follows a particle in the bulk material. This charge may be adequate to induce an SCR (silicon controlled rectifier) turn-on effect in the device's substrate (see fig. 5) thereby inducing a low impedance path between power and ground. Single Event Burnout, like SEL, is quite often destructive, however it is restricted to power HEXFETS that arc relativel y rare in computer systems.

All of these single-event failure modes arc caused by energetic charged particles from the sun, the galactic cosmic rays background, or intense magnetospheres (e.g. Jupiter). These particles range in size from protons to heavy atomic nuclei like ionized oxygen.

Energetic particle flux is a strong function of intermittent solar output that can fluctuate by many orders of magnitude during large solar flare events [22]. In fact, during one solar flare event in 1989, during one week in October, the total flux of protons with energies greater than 10 McV exceeded the total for the entire previous 11-year solar cycle (had

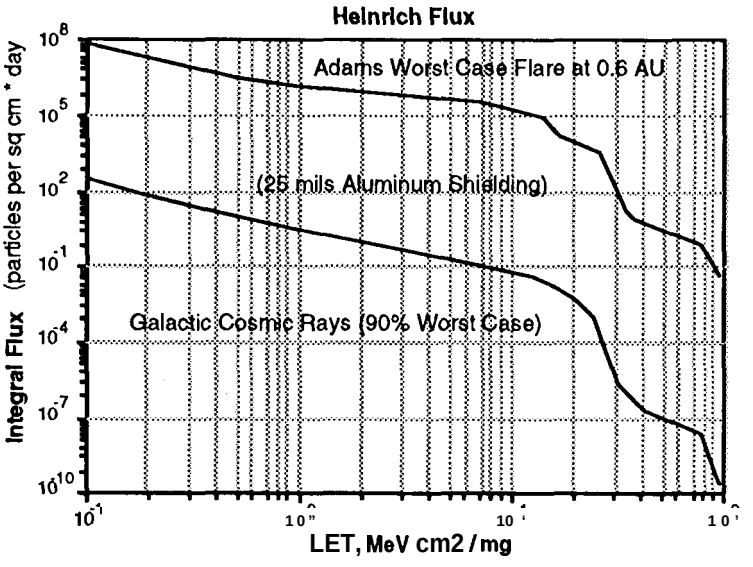

Figure 4

astronauts been on the moon during this time, they would have received a lethal radiation dose, about 500rad, in only two days) [24].

For the bulk of most missions, solar activity remains very low compared with the galactic cosmic ray (GCR) background, therefore it is convenient to specify maximum solar flux output separately from the relatively constant GCR particle flux.

Similar to TID radiation, the SEU (SEL) LET threshold for a component is typically expressed in units of LET (Linear Energy Transfer) and is the maximum amount of energy per unit length (dE/dx) required to be deposited on a sensitive region of the component (typically normalized to silicon) before bit upsets (or latchup) may occur. Below that threshold, the failure phenomena dots not occur because the particle does not have enough energy to cause state changes.

LET is specified in MeV cm^2 / mg. Although the units are different, numbers expressed in these units can be translated to charge per unit length by dividing by the mass density of silicon and translating MeV into number of charge carriers by dividing by 3.6 eV per electron-hole pair. Particles that can cause SEU or SEL are also specified in LET. A plot of the number of particles in a particular environment greater than or equal to a particular LET value is called the Heinrich Integral Flux (or just Heinrich Flux).

Fig. 4 shows the integral flux for two environments, GCR and 0.6 AU solar flare. These are not necessarily representative of all particle environments, but it serves to illustrate that the number of particles with energies above a given LET tend to drop off exponentially with increasing LET. Therefore small increases in a component's LET threshold greatly reduces the number of particles that are of concern. From the figure, it can be seen that there are nearly 10^8 particles that can deposit at least 0.1 LET incident on a 1 cm^2 silicon device every day in a 10% worst case solar flare at 0.6 AU. If the device was uniformly distributed with very many identical sensitive devices (memory bits) that covered the entire chip surface and those memory bit's LET threshold was 0.1 or less, then one would expect 10^8 bit upsets per day! Fortunately things aren't all that bad, even the "softest" digital devices (e.g. high density dynamic memories) have SEU LET thresholds well above 1. Another saving grace is that total area of the chip that contains sensitive volumes, also known as the device cross section, is usually only a small fraction of the chip's total area. Even so, under worst case solar flare conditions the SEU rate can be as high as 100's of bit upsets per chip-day and as bad as 10's of bit upsets per chip-day in the GCR environment [25] [8].

For processor chips there is a slight additional saving. It has been shown via simulation [26] for some CISC processors that about 40% of the bit upsets occurring in the processor actually induce operational errors (many register bits as at various times "don't care" bits). This may not remain true in highly optimized piplined RISC machines.

For control applications where computer components play a key role in maintaining spacecraft integrity, high enough SEU-induced upsets in non-EDAC protected data storage may preclude the flight computer from doing anything other than performing constant SEU error detection and recovery [27] [28]. General] y speaking, if SEU error detection and recovery in the flight computer occupy any more than a small fraction (<<10%) of the available processor throughput, then it is highly unlikely that any method other than SEU hardening or concurrent hardware error detection and correction will suffice. Both of these methods may preclude the use of off-the-shelf flight computer technology if the expected SEU rate is very high.

This situation is exacerbated if the device shows susceptibility to SEL. The SCR effect common in non-epitaxial bulk CMOS (see fig. 5), can greatly limit the mean-time-to-failure of these components. As soon as a particle with high LET deposits enough charge to forward bias the PN and NP junctions in the device, the parasitic SCR circuit will turn on and provide a low impedance current path from power to ground. If the current density is high enough, the device will be destroyed. For this mason, typical SEL LET threshold requirements range from 40 to as high as 110 for long duration missions.

If it can be shown that the latchup current density is not so high as to inflict damage, temporarily removing power from the device will cause the SCR effect to cease. Some low criticality spacecraft applications [29] have placed latchup current detectors external to the susceptible components. On latchup detection, the power to the whole computer is cycled off then on (computer reset is not adequate to eliminate this effect).

To avoid the expense and risk associated with high coverage software (and hardware) SEU and SEL recovery methods, commercial device SEE hardening programs can be attempted. This can and has been done a number of times at greater expense than for TID hardness process changes [16]. These costs are high because SEU hardness can not always be achieved with simple process modifications. Two approaches predominate: electrical modification of the memory cell and redesign of the chip for use in harder processes (e.g. CMOS SOS or SOI processes from bulk CMOS). If the original design of the chips was done using a high level portable logic library (standard cell or gate array), then "porting" the design to another logic library expressly designed for SEU hardness in the electrical design of the memory elements (e.g. larger gate capacitance, cross-strap feedback resistors) can be done at reasonable expense. The other approach entails transferring the masks of the original (bulk) CMOS design to a foundry that specializes in SEU hard CMOS-SOS or CMOS-SOI. These processes yield higher LET threshold components because the size of the sensitive volume is much smaller than for bulk CMOS. Mask translations from bulk to S01 or SOS can be difficult depending on the geometrical design rules used in the original component [30] [31]. Both of these approaches rely on the original chip vendor to be cooperative and be willing to give or sell the details of the design to the spacecraft system designers.

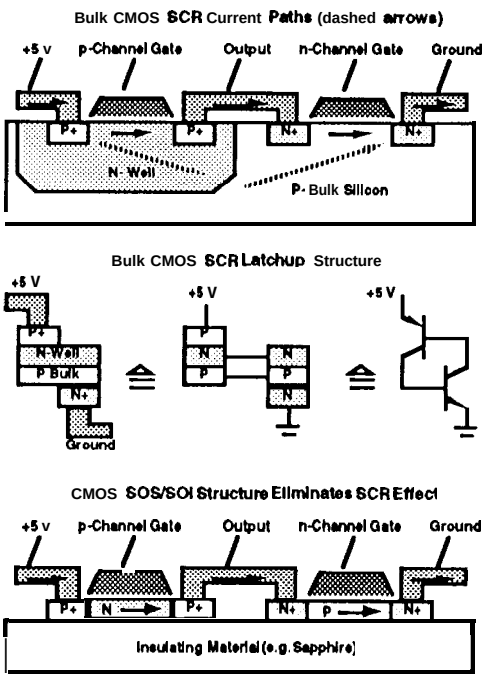

Figure 5

With the advent of portable design specifications and portability tools (e.g. VHDL, Synopsis) in the future, the former approach is likely to become much cheaper, and in fact,

portable application specific designs are already becoming more common place [23].

SEU and SEL testing requires high energy, high-Z particle accelerators in order to provide adequate simulation of the cosmic ray environment. As a result, there arc very few of these SEU test facilities in existence. In fact the number of these facilities may be shrinking, one of the more popular, the Bevatron accelerator at the Lawrence Berkeley Laboratory, is closing down. This leaves the Single Event Upset Test Facility at the Brookhaven National Laboratory as onc of the accessible few. In the future, SEU and SEL rate determination will be based more on computer modeling approaches [32].

In thc near future, as the minimum feature size of new computer components shrinks below 1 μm, the critical charge will drop, thereby making the SEU (and SEL) susceptibility of these components higher [33], even in terrestrial applications. Fortunately, this is partially offset due to a corresponding reduction in the device cross section. Hopefully, the terrestrial SEU phenomena will cause chip manufactures to apply hardening (and or fault tolerance) techniques to all computer designs, not just spacecraft computers.

Fault Tolerance for Cost Saving

If the SEL rate is acceptable and the SEU rate is not overwhelming, hardware and software error detection and recovery mechanisms (i.e., fault tolerance) can go a long way in making an SEU "soft" flight control computer appear to be much harder. However, these techniques come at a cost of additional complexity, mass, power, volume and performance reduction. For example, error detection and correction codes can be easily added between the processor and its memory to scrub out single bit memory errors. This adds to the processor mass and power, but it can also adversely affect processor performance. The processor-memory interface is especially critical to RISC machines that require low latency memory accesses to be competitive with CISC architectures.

Basing spacecraft computers on standard-off-the-shelf computers may limit the ability to add fault tolerance as these computers arc typically not designed for high concurrent error detection techniques (such as register and bus parity checkers, lock-stepped self checking, hardware roll-back/retry, etc.) [34] [35].

Other high level fault tolerance techniques may be applied to circumvent these limitations. A fully redundant pair of synchronized computers operating in parallel with output comparison hardware can provide very high error detection coverage mechanism for reasonable hardware costs. Hardware costs can be reduced further if synchronization and output voting is performed at the software task level [36]. Triple or quad-redundancy can allow "operate-through" capability if implemented carefully.

Design for Quality

Standard military quality standards (e.g. Mil-M-385 10, and MIL-I-38535) not just raised the quality of individual components, they have also made quality specification easier. The two main quality levels commonly used for space applications 38510 (JAN) class B (lower quality) and class S (higher quality) arc both arc independent of required radiation levels (sw above).

If the purposeofMil-M-38510 is to screen out lemons, then the purpose MIL-I-38535 is to ensure that manufacturers make no lemons. MIL-I-38535 was invented to expedite the process of getting ncw chip designs qualified (any ncw chip design, even ASICs, if followed through the approved process will be certified) hence lower the cost. The associated parts certification levels arc class Q (lower quality) or class V (higher quality).

It is not clear which approach will ultimately provide lower cost to the spacecraft designer. However, components procured to these standards can be assured to be at least a factor of 10 higher in recurring costs than if no quality level was specified. The option

exists for a spacecraft computer system designer to rely only cm the intrinsic quality level that most manufactures maintain without specifying component class (other than temperature and radiation requirements). If many small spacecraft arc built then the individual failure rate may not be significantly different if class S or B parts are used.

SUMMARY

The space of options available to spacecraft control electronic systems designers has certainly widened in recent years. The rate of change of terrestrial technology has offered hope that small, low cost spacecraft will directly benefit. However the traditional environmental concerns have not disappeared and arc not likely to go away any time soon. Total Dose and Single Event Effects (SEL in particular) still make radiation hardening the easiest method of getting there from here. For SEU (and SEL) rates that are not intolerable, small, high performance computer technologies enables the use of task level software implemented fault tolerance as an approach to reducing functional failure rate. If SEL effects arc convincingly non-destructive, latchup detection and power cycling circuitry can circumvent the SCR cffeet. As feature sizes shrink, total dose hardness due to trapped charges will improve. Also as feature sizes shrink, SEU and SEL effects will become more common in non-flight computers, hence manufacturers will have to apply hardening and/or fault tolerance techniques to terrestrial computers.

As long as space-unique effects remain, flight computer costs will remain higher than comparable terrestrial computers, however, costs will go down.

ACKNOWLEDGEMENTS

The author wishes to thank Robert Rasmussen for his patient and thoughtful discussions on this complex subject. This work was carried out by the Jet Propulsion Laboratory, California Institute of Technology under a contract with the National Aeronautics and Space Administration.

REFERENCES

1. K. Ports and C. Tabbert, "Commercial Applications Benefit from Dual-use Technology", Electronic Engineering Times, Nov. 9, 1992, p. 62
2. GAO Report to the Subcommittee on Science, Space, and Technology, "NASA Efforts to Develop and Deploy Advanced Spacecraft Computers", Mar. 1989
3. 1.. Holcomb, "Overview of NASA's On-board Computing Technology Program", Digest of Papers, IEEE Compcon Spring 80, p 117
4. D. F. Barbe, "VHSIC Systems and Technology", IEEE Computer, Vol. 14, No. 2, Feb. 1981, p. 13
5. G, Messenger, M. Ash, "The Effects of Radiation on Electronic Systems", Van Nostrand Rienhold Company 1986
6. R. Rasmussen, "Spacecraft System Electronics Design for Radiation Tolerance", Proc. of the IEEE, vol. 76 no. 11, Nov. 1988
7. J. Greenbaum, "Radiation Effects - The 1 Iows and Whys", IEEE Circuits and Systems Magazine, vol. 1, no. 4, p. 11
8. P. Lewkowicz and L. Richter, "Single Event Upsets in Spacecraft Digital Systems", Proc. International Telemetering Conference, Oct. 1984, p. 781
9. J. Bradford and P. Vail, "Cosmic Ray Effects in Spaceborne Electronics", Proc. of Digital Avionics Systems Conf, Nov. 1979, p. 81
10. 1 I. Curnow and B. Wickman, "A Synthetic Benchmark", The Computer Journal, Vol. 19 No. 1, Feb. 1976, p. 43
11. D. A. Patterson, "Reduced Instruction Set Computers", Comm. ACM 281985, 8-21
12. J. Hennessy, et al, "I Iardware/Software Tradeoffs for Increased Performance", Proc. S ymposium on Architectural Support for Programming Languages and Operating Systems, 1982
13. J. Schitt, R. Creasey, F. Gomez-Molinero, "The Development of High Performance Flight Computers

for the European Space Programmes", Proc. 40th Congress of the international Astronautical Federation, IAF-89-045.

14. '1', Naegal, "The Rise and Fall and Rise Again of Mil-Std1 750A", Military and Aerospace Electronics, Sept. 1990, p 19
15. S.C. Lee and J, R. Hayes, "Development of a Forth Language Directed Processor using Very Large Scale Integrated Circuitry", John Hopkins APL Technical Digest, Vol 10, No. 3, 1989, p. 216
16. A. Giddings et. al., "SEU Immune ICs for Project Galileo", IEEE Transactions on Nuclear Science, 1985
17. M. E. Thimlar, W.E. Crane, D.R. O'Bell, "Furure Space-Based Computer Processors", AIAA Aerospace America, Mar. 1984 p.78
18. P.C. Carney, "Selecting On-Board Satellite Computer Systems", IEEE Computer, Apr. 1983, p. 35
19. D. J. Theis, "Spacecraft Computers: State-of-the-Art Survey", IEEE Computer, Apr. 1983, p. 85
20. J. Burr and A. Peterson, "Ultra Low Power CMOS Technology", 3rd NASA Symp. on VLSI Design, 1991, p. 11.1.1
21. R. Panwar and L, Alkalaj, "Low Energy CMOS for Space Applications", 4th NASA S ymp. on VLSI Design, 1992, p. 10.4,1
22. 1,. David, "CRRES Finds Radiation Tough on Tiny Electronics", Space News, June 22, 1992
23. M.E. Fracman et al (API.), "Experience with Custom processors in Space Flight Applications", 3rd NASA Symp. on VLSI Design, 1991, p. 1.1.1
24. News Notes, "Devastating Solar Activity in 1989", Sky and Telescope, June, 1990
25. J. Adams et. al., "Cosmic Ray Effects on Microelectronics" Part I, NRL Report No. 4506, Aug. 1981; Part 11, No. 5099, May 1983; Part IV, NRL Report No. 5901, Dec. 1986
26. G, Burdick, E Kopf, D. Meyer, "The Galileo Single Event Upset Solution and Risk Assessment", Pmt. of AAS Guidance and Control Conference, Feb. 1986, p.199
27. R, Rasmussen, "Computing in the Presence of Soft Bit Errors", Proc. of American Controls Conference, June 1984
28. S. Cummingham, J. Banasiak, and C von Flotow, "Living with things that go bump in the night", Proc. of AAS Guidance and Control Conference, Feb. 1985, p.469
29. R, Maurer, J. Kinnison, "Single Event Upset and Latchup Sensitive Devices in Satellite Systems", Proc. AIAA/USU Conf. on Small Satellites, Sept. 21, 1992
30. J. Colinge, "Problems and Issues in SOI CMOS Technology", Proc. IEEE International SOI Conference, 1991, p. 126
31. B. Offord, "l~ully-Depleted Silicon-on-Sapphire and Its Application to Advanced VLSI Design", 4th NASA Symp. on VLSI Design, 1992, p. 4.3.1
32. J. Kinnison, "Single Event Phenomena: Testing and Prediction", 4th NASA Symp. on VLSI Design, 1992, p. 1.2.1
33. E. Burke, "The Impact of Component Technology Trends on the Performance and Radiation Vulnerability y of Microelectronic Systems", IEEE Conf. on Nuclear and Space Radiation Effects Short Course, July 1985
34. D. Sicwiorck and R. Swam, "The Theory and Practice of Reliable System Design", Digital Press 1982
35. D. Rennels, "Fault Tolerant Computing - Concepts and Examples", IEEE Transactions on Computers, C-33, Dec. 1984, p. 1116
36. M. Malhotra, "Parallel Processing On-Board the UoSAT Microsatellite", Proc. AIAA/USU Conf. on Small Satellites, Sept. 21, 1992